

REMARKS/ARGUMENTS

Claims 1-14 and 16-19 are pending in the application and have been rejected. Claim 18 has been amended. Applicant respectfully requests reconsideration.

REJECTIONS UNDER 35 U.S.C. §102.

The Office Action has rejected claims 1-6, 9-13 and 9-19 under 35 U.S.C. §102 as being anticipated by US Patent No. 6,415,359 issued to Kimura et al. (hereafter, "Kimura"). Applicant respectfully traverses the rejection for the following reasons.

Kimura does not teach or suggest using non-volatile secondary storage. The Office Action cites a disk cache for this limitation. See Office Action at page 3 (citing col. 3, line 55 to col. 4, line 15). A disk cache is defined: "A portion of a computer's random access memory (RAM) set aside for temporarily holding information read from disk..." *Microsoft Computer Dictionary, Fifth Edition*, Microsoft Press (2002). RAM is defined as follows: "Acronym for random access memory. Semiconductor-based memory that can be read and written by the central processing unit (CPU) or other hardware devices... The term RAM, however, is generally understood to refer to **volatile** memory that can be written to as well as read." (emphasis added) *Id.*

Claim 1 requires storing only *strategically selected storage data* in the second level of storage based on energy-conserving criteria. Kimura neither teaches nor suggests this feature. The Office Action cites col. 4, lines 16-25 and the Abstract of Kimura for this element. However that part of Kimura does not discuss any strategically selected storage data.

The Office Action alleges that the foregoing argument is a general allegation that the claims define a patentable invention. Applicant respectfully reminds the Examiner that the Examiner (or the Board, if the Board is the first body to raise a particular ground for rejection) "bears the initial burden . . . of presenting a *prima facie* case of unpatentability." *In re Oetiker*, 977 F.2d 1443, 1445, 24 USPQ2d 1443, 1444 (Fed. Cir. 1992). Applicant's argument is a specific argument that the Examiner failed to carry the burden of showing a *prima facie* case by pointing out that the cited part of the prior art did not discuss the specific limitation of storing

only strategically selected data. In fact that is a correct statement. The Office Action misconstrues the claim term “strategically selected” to mean “specific” yet cites no dictionary, part of the specification, or other authority for this. See Office Action at page 13. If a cited reference fails to mention a claim limitation an applicant can and should assert that the limitation is missing. That is not a **general** allegation that the claims define a patentable invention.

Claims 2-15 are either directly or indirectly dependent on claim 1 and are patentable for at least the same reasons that claim 1 is patentable.

Claim 16 recites first and second levels of non-volatile memory and a controller for storing only strategically selected storage data in the second level of storage. As noted above Kimura does not teach or suggest first and second levels of non-volatile storage or storing only strategically selected storage data. Therefore claim 16 is not anticipated by Kimura.

Claim 17 relates to a method for managing storage of information and includes two levels of non-volatile storage and a step of storing only strategically selected storage data in managed storage. As noted above Kimura does not teach or suggest two levels of non-volatile storage or storing only strategically selected storage data. Therefore claim 17 is not anticipated by Kimura.

Claim 18 relates to computer readable medium and includes elements of storing only strategically selected storage data in managed storage and storing all storage data in non-managed non-volatile storage when the operating state of the system does not satisfy the one or more energy-conserving criteria. For reasons discussed with respect to claim 1, Kimura does not teach or suggest these limitations. Therefore claim 18 is not anticipated by Kimura.

Claim 19 recites an information handling system that includes two levels of non-volatile storage and an arbiter for storing only strategically selected storage data in second level storage. For reasons discussed above, Kimura does not teach or suggest these limitations. Therefore claim 19 is not anticipated by Kimura.

REJECTIONS UNDER 35 U.S.C. §103.

The Office Action rejected claims 7 and 8 as unpatentable over Kimura in view of Thelander (U. S. Patent Application 2003/0009705). The Office Action concedes that Kimura is silent on "system stores current user profiles and the system state information comprises whether storage input/output data are associated with a current user profile." However, the Office Action contends that Thelander teaches this element and that it would have been obvious to one skilled in the art to modify Kimura according to Thelander such that the limitations of claim 7 are met. The motivation cited for making the combination is that "the power magnitude profile may include multiple power settings or power schemes with the same schedule, so that the user may select between different power settings or schemes to be implemented." Page 10 of the Office Action. The Office Action cites paragraph 48 of Thelander for this motivation. Paragraph 48 reads:

"With some embodiments of the invention, the power management profile may include multiple power settings or power schemes with the same schedule, so that the user may select between different power settings or schemes to be implemented during the scheduled time period. In these embodiments, the client unit 303 may cause the operating system to implement a default power setting or scheme for the scheduled period according to any desired criteria, such as having a preselected default setting or scheme, or using the setting or scheme most recently selected by the user. Also, with some embodiments of the invention, the power management profile may allow any user (or one or more specified users) to modify or override parameters of the profile. With these embodiments, the interface 401 may, for example, permit the user to modify the settings in fields 445 and/or 447 and 449, or provide an override button that allows the user to override implementation of the power management profile."

As the quoted section reveals, Thelander concerns a power management profile and not a user profile. Thelander does not suggest the second part of claim 7 either. There is no discussion in Thelander of the system state information comprising storage input/output data associated with a current user profile. Therefore, Thelander does not suggest or motivate the modification of Kimura as suggested by the Office Action.

With respect to claim 8, the Office Action also contends that "the combination of Kimura and Thelander teach, wherein the system stores current user preferences and the system state information comprises whether storage input/output data are associated with current user

preferences.” The above-quoted part of Thelander is cited in support of this finding. However, neither Kimura nor Thelander teach, suggest or motivate the caching of only strategic information for power savings.

Therefore, Applicant respectfully submits that claims 7 or 8 would not have been obvious in view of the combination of Kimura and Thelander. The concepts of “first” and “storing only strategically selected information” is not taught or suggested in the combination of references. Thelander contains no teaching or motivation to modify Kimura as suggested by the Examiner.


Claim 11 has been rejected under 35 USC §103 as being obvious over Kimura in view of Applicant's admitted prior art. The Office Action concedes that Kimura does not disclose the use of Flash memory for the second levels of storage hierarchy. Claim 11 is dependent on claim 1 and is patentable for the same reasons as claim 1. Applicant has not admitted that the claimed combination is old. The statement made in the Background is “These media *could be* alternate non-volatile memory such as Flash memory, ...(emphasis added)”. That is a statement of possible solutions not an admission that those devices exist. The mere existence of Flash memories does not establish the obviousness of using the Flash memories in the combination of claim 8. The claimed subject matter must be considered as a whole. 35 USC §103. This rejection in effect attempts to show that one skilled in the art might find it obvious to try various combinations. However, that is not the standard of 35 U.S.C. §103. *In re Geiger*, 815 F.2d 868, 2 USPQ2d 1276 (Fed. Cir. 1987). Obvious to try is not the test for obviousness. *Id.* Such suggestion or motivation to combine cannot come exclusively from the disclosure of the invention undergoing examination.

Claim 12 has been rejected under 35 USC §103 as being obvious over Kimura in view of Atkinson (US Patent 6,029,249). Claim 12 is indirectly dependent on claim 1 and is patentable for at least the same reasons discussed herein with respect to claim 1.

Serial No. 10/674,926
Filed: 09/30/2003
AMENDMENT

For the foregoing reasons, Applicant respectfully requests allowance of the pending claims and that a timely Notice of Allowance be issued in this case.

Respectfully submitted,


Michael J. Buchenhorner
Reg. No. 33,162

Electronically filed on Date: March 27, 2007

Attachment: two photocopied pages from *Microsoft Computer Dictionary, Fifth Edition*,
Microsoft Press (2002)

disk cache

By a disk. Because disk devices are slow compared with the CPU, it is not efficient to access the disk for only one or two bytes of data. Instead, during a read, a large chunk of data is read and stored in the disk buffer. When the program wants information, it is copied from the buffer. Many requests for data can be satisfied by a single disk access. The same technique can be applied to disk writes. When the program has information to store, it writes it into the disk buffer area in memory. When the buffer has been filled, the entire contents of the buffer are written to the disk in a single operation.

disk cache *n.* A portion of a computer's random access memory (RAM) set aside for temporarily holding information read from disk. A disk cache does not hold entire files, as does a RAM disk (a portion of memory that acts as if it were a disk drive). Instead, a disk cache is used to hold information that either has recently been requested from disk or has previously been written to disk. If the required information remains in a disk cache, access time is considerably faster than if the program must wait for the disk drive mechanism to fetch the information from disk. *See also* cache. *Compare* disk buffer.

disk cartridge *n.* A removable disk enclosed in a protective case. A disk cartridge can be used by certain types of hard disk drives and related devices, such as the external data storage units known as Bernoulli boxes.

disk controller *n.* A special-purpose chip and associated circuitry that directs and controls reading from and writing to a computer's disk drive. A disk controller handles such tasks as positioning the read/write head, mediating between the drive and the microprocessor, and controlling the transfer of information to and from memory. Disk controllers are used with floppy disk drives and hard disks and can either be built into the system or be part of a card that plugs into an expansion slot.

disk copy *n.* The process of duplicating a source disk's data and the data's organizational structure onto a target disk. *See also* backup.

disk crash *n.* The failure of a disk drive. *See also* crash¹.

disk directory *n.* An index of the files on a disk, analogous to a card catalog. A disk directory includes information about the files, such as their names, sizes, dates of creation, and physical locations on the disk. *See also* directory.

disk drive *n.* An electromechanical device that reads from and writes to disks. The main components of a disk

disk mirroring

drive include a spindle on which the disk is mounted, a drive motor that spins the disk when the drive is in operation, one or more read/write heads, a second motor that positions the read/write heads over the disk, and controller circuitry that synchronizes read/write activities and transfers information to and from the computer. Two types of disk drives are in common use: floppy disk drives and hard disk drives. Floppy disk drives are designed to accept removable disks in either 5.25-inch or 3.5-inch format; hard disk drives are faster, high-capacity storage units that are completely enclosed in a protective case.

disk driver *n.* A device driver that is added to a system to support a specific manufacturer's disk device. *See also* device driver.

disk duplexing *n.* *See* disk mirroring.

disk envelope *n.* The paper container that holds a 5.25-inch floppy disk and its attached jacket. The disk envelope protects exposed surfaces of the disk from dust and other foreign material that can scratch and otherwise damage the surface, resulting in the loss of recorded data. *See also* disk jacket.

diskette *n.* *See* floppy disk.

disk farm *n.* A number of disk drives in a single location used together to store or process vast quantities of information, such as scientific data, years' worth of corporate sales figures, large numbers of graphic images, or telephone company billing records. Current disk farms consist of magnetic or optical disks and can hold terabytes of information. In older usage, disk farms were sometimes known as "Laundromats" because they contained large drives referred to in jargon as "washing machines." *See also* server farm.

disk interface *n.* 1. The circuitry that connects a disk drive to a computer system. 2. A standard for connecting disk drives and computers. For example, the ST506 standard for connecting hard disks to computers is a disk interface standard.

disk jacket *n.* The protective plastic sheath that covers a floppy disk.

diskless workstation *n.* A station on a computer network that is not equipped with a disk drive and that uses files stored in a file server. *See also* file server.

disk memory *n.* *See* virtual memory.

disk mirroring *n.* A technique in which all or part of a hard disk is duplicated onto one or more other hard disks,

D

entication Dial-In
net protocol in
s authorization and
k server to which a
entication, commu-
).

—for example, 2 in
/stem, 8 in the octal
tem. See also base

system for represent-
possible digits
s for them, the num-
subtracting each of
e largest possible
example, in a sys-
radix-minus-1
ng any number to its
the largest possible
-, 99,999). Adding
our example,
ve digits are used, the
ny number in the sys-
plus 1, because $-a +$
minus-1 complement
sily formed electroni-

character that separates
the fractional port-
point is the decimal

g algorithm that sort
successive parts of
ng a list of numbers
ed by the hundreds
in each list, one at a
0 lists based on the
lists is sorted by the
most efficient whe-
ies, which simplifies
(?) and reduces the
at most two lists).
ve asymmetric digi-
ed version of ADSL

(asymmetric digital subscriber line) that is capable of adjusting transmission speed (bandwidth) based on signal quality and length of the transmission line. As the signal quality improves or deteriorates while a transmission line is being used, the transmission speed is adjusted accordingly. See also ADSL, xDSL.

rag *n.* Irregularity along the left or right edge of a set of lines of text on a printed page. Rag complements justification, in which one or both edges of the text form a straight vertical line. See the illustration. See also justify, ragged left, ragged right.

Ragged right	Justified	Ragged left
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____

Rag.

ragged left *adj.* Of, relating to, or being lines of text whose left ends are not vertically aligned but form an irregular edge. Text may be right-justified and have a ragged left margin. Ragged-left text is used infrequently—typically, for visual effect in advertisements. See also rag, right-justify.

ragged right *adj.* Of, relating to, or being lines of text whose right ends are not vertically aligned but form an irregular edge. Letters and other word-processed documents are commonly left-justified, with ragged-right margins. See also left-justify, rag.

RAID *n.* Acronym for redundant array of independent (or inexpensive) disks. A data storage method in which data is distributed across a group of computer disk drives that function as a single storage unit. All the information stored on each of the disks is duplicated on other disks in the array. This redundancy ensures that no information will be lost if one of the disks fails. RAID is generally used on network servers where data accessibility is critical and fault tolerance is required. There are various defined levels of RAID, each offering differing trade-offs among access speed, reliability, and cost. See also disk controller, error-correction coding, Hamming code, hard disk, parity bit, server (definition 1).

RAID array *n.* See RAID.

RAM *n.* Acronym for random access memory. Semiconductor-based memory that can be read and written by the central processing unit (CPU) or other hardware devices. The storage locations can be accessed in any order. Note that the various types of ROM memory are capable of random access but cannot be written to. The term RAM, however, is generally understood to refer to volatile memory that can be written to as well as read. Compare core, EPROM, flash memory, PROM, ROM (definition 2).

RAMAC *n.* 1. Acronym for Random Access Method of Accounting Control. Developed by an IBM team led by Reynold B. Johnson, RAMAC was the first computer disk drive. It was introduced in 1956. The original RAMAC consisted of a stack of 50 24-inch platters, with a storage capacity of 5 megabytes and an average access time of 1 second. 2. A high-speed, high-capacity disk storage system introduced by IBM in 1994. Based on the original RAMAC storage device, it was designed to fulfill enterprise requirements for efficient and fault-tolerant storage.

Rambus DRAM *n.* See RDRAM.

Rambus dynamic random access memory *n.* See RDRAM.

RAM cache *n.* Short for random access memory cache. Cache memory that is used by the system to store and retrieve data from the RAM. Frequently accessed segments of data may be stored in the cache for quicker access compared with secondary storage devices such as disks. See also cache, RAM.

RAM card *n.* Short for random access memory card. An add-in circuit board containing RAM memory and the interface logic necessary to decode memory addresses.

RAM cartridge *n.* See memory cartridge.

RAM chip *n.* Short for random access memory chip. A semiconductor storage device. RAM chips can be either dynamic or static memory. See also dynamic RAM, RAM, static RAM.

RAM compression *n.* Short for random access memory compression. This technology was an attempt by a number of software vendors to solve the problem of running out of global memory under Windows 3.x. Compression of the usual contents of RAM may lessen the system's need to read or write to virtual (hard disk-based) memory and thus speed up the system, as virtual memory is much slower than physical RAM. Because of the falling prices of RAM and the introduction of operating systems that handle RAM

R